arXiv:1611.07096v1 [stat.ML] 21 Nov 2016

Structured Prediction by Least Squares Estimated Conditional Risk Minimization

Chong Yang Goh
Operations Research Center
Massachusetts Institute of Technology
cygoh@mit.edu

Patrick Jaillet
Department of EECS
Massachusetts Institute of Technology
jaillet@mit.edu

Abstract

We propose a general approach for supervised learning with structured output spaces, such as combinatorial and polyhedral sets, that is based on minimizing estimated conditional risk functions. Given a loss function defined over pairs of output labels, we first estimate the conditional risk function by solving a (possibly infinite) collection of regularized least squares problems. A prediction is made by solving an auxiliary optimization problem that minimizes the estimated conditional risk function over the output space. We apply this method to a class of problems with discrete combinatorial outputs and additive pairwise losses, and show that the auxiliary problem can be solved efficiently by exact linear programming relaxations in several important cases, including variants of hierarchical multilabel classification and multilabel ranking problems. We demonstrate how the same approach can also be extended to vector regression problems with convex constraints and losses. Evaluations of this approach on hierarchical multilabel classification show that it compares favorably with several existing methods in terms of predictive accuracy, and has computational advantages over them when applied to large hierarchies.

1 Introduction

Many important tasks in machine learning involve predicting output labels that must satisfy certain joint constraints. These constraints help restrict the search space and incorporate domain knowledge about the task at hand. For example, in hierarchical multi-label classification, the goal is to predict, given an input, a set of labels that satisfy hierarchical constraints imposed by a known taxonomy. More generally, the output space can also depend on the observed inputs, which is commonly the case in language, vision and speech applications. Collectively, these problems are broadly referred to as *structured prediction*.

Formally, the goal of structured prediction is to learn a predictor $h : \mathcal{X} \mapsto \mathcal{Y}$ that maps an input $x \in \mathcal{X}$ to some y in a structured output space $\mathcal{Y}$. Unlike the classical problems of binary classification ($\mathcal{Y} = \{-1, 1\}$) and regression ($\mathcal{Y} = \mathbb{R}$), a defining challenge in structured prediction is that $\mathcal{Y}$ can be a highly complex space, usually representing combinatorial structures like trees, matchings, and vertex label assignments on graphs, or any arbitrary set of vectors in a real vector space. The difficulty of the task naturally depends on the geometry of $\mathcal{Y}$ and the choice of loss function $\ell : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}_+$.

Many existing structured prediction methods can be understood as learning a scoring function $F(y, x)$, which assigns to each $y \in \mathcal{Y}$ its compatibility score with some input x. A prediction is made by solving an auxiliary optimization problem that finds some $y^* \in \arg\max_{y \in \mathcal{Y}} F(y, x)$. Conditional random fields (CRFs) [1], Max-Margin Markov Networks (M^{3}N) [2] and Structured Support Vector Machines (SSVM) [3] are three commonly used methods that fit into this framework. CRFs directly model the conditional distribution $F(y, x) := p_{Y|X}(y|x)$ with a graphical model, thus maximizing the compatibility score is equivalent to maximum a posteriori estimation. In M^{3}N and SSVM, $F(y, x)$

Submitted to 29th Conference on Neural Information Processing Systems (NIPS 2016) on May 20, 2016.

is expressed in a linear (or log linear) form $\langle w, \phi(x, y)\rangle$, where $\phi(x, y)$ is a joint feature representation of the input-output pair. The weight vector vector w is found by solving a generalized max-margin problem that maximizes separation of the true labels from others by their score differences, which are usually augmented with a loss function.

In this paper, we propose an approach to structured prediction by *least squares estimated conditional risk minimization*, which can also be characterized as one that learns a scoring function $F(y, x)$. Specifically, the scoring function (to be minimized over) can be interpreted as a direct estimate of the *conditional risk function*, $R(y|x) := \mathbb{E}_Y[\ell(y, Y)|x]$, where ℓ is a loss function of choice. Notably, like M^{3}N and SSVM, our approach can also use kernels to learn from a rich class of hypotheses, but avoids the nontrivial *pre-image problem* of the former, which involves finding some y that maximizes the inner product $\langle w, \phi(x, y)\rangle$. Our main contributions are summarized as follows: (i) we derive a closed form expression for the estimated conditional risk function, in the form of a weighted sum of loss functions, as solutions of a collection of regularized least squares problems (Section 2) (ii) we show that in a certain class of problems with discrete combinatorial outputs and additive pairwise losses, the auxiliary problem can be solved efficiently by exact linear programming relaxation (Section 3) (iii) we extend this approach to convex output spaces and learning additive models of conditional risks (Section 4) (iv) we derive an excess risk bound for the algorithm under a parametric class of $R(y|x)$ and finite $\mathcal{Y}$, and show that it scales logarithmically with $|\mathcal{Y}|$ (Section 5). When applied to hierarchical multilabel classification (Section 6), our approach achieves predictive accuracies that are competitive with several existing algorithms, with added computational advantages that allow it to scale well with the hierarchy size.

2 Least Squares Estimated Conditional Risk Minimization (LS-ECRM)

2.1 Preliminaries

Throughout this paper, we represent $\mathcal{X} \subseteq \mathbb{R}^p$ and $\mathcal{Y} \subseteq \mathbb{R}^d$ as collections of vectors in a real vector space. Depending on the application, $\mathcal{Y}$ can be discrete (e.g., a combinatorial subset of $\{0, 1\}^d$) or otherwise (e.g. a polytope). As in most supervised learning settings, we have a training set, $S = [(x^{(i)}, y^{(i)}) \in \mathcal{X} \times \mathcal{Y} : i = 1, \ldots, m]$, where the samples are i.i.d. and drawn from some fixed distribution $\mathbb{P}_{X,Y}$. We define $R(y|x) := \mathbb{E}_Y[\ell(y, Y)|x]$ to be the conditional risk of predicting label y having observed an input x, and use $\hat{R}$ to denote its estimate. Here the conditional expectation is defined with respect to $\mathbb{P}_{Y|X}$.

In summary, the proposed method can be described in two steps:

1. **Training**: Learn a set of functions indexed by y, $\{\hat{R}(y|\cdot) : y \in \mathcal{Y}\}$ by solving a (possibly infinite) collection of regularized least squares problems.
2. **Prediction**: Given an input x, predict $h(x) = y^*$ by solving an auxiliary optimization problem,
$$y* \in \arg\min_{y \in \mathcal{Y}} \hat{R}(y|x). \tag{1}$$

2.2 Training

We begin with two basic observations: for any fixed $y \in \mathcal{Y}$, (i) the random loss function $\ell(y, Y)$ can be written as $\ell(y, Y) = R(y|X) + \varepsilon(y, X)$, where $\varepsilon(y, X)$ is a zero-mean random variable that depends (randomly) on X and (deterministically) on y (ii) $R(y|\cdot) \in \arg\min_f \mathbb{E}[(f(X) - \ell(y, Y))^2]$. This decomposition suggests a regression approach to estimating $R(y|\cdot)$: for every $y \in \mathcal{Y}$, we posit that $R(y|\cdot)$ lies in some function space $\mathcal{H} \subseteq \{f : \mathcal{X} \mapsto \mathbb{R}\}$ and estimate it by solving a regularized least squares problem,[1]

$$\hat{R}(y|\cdot) \in \arg\min_{f \in \mathcal{H}} \sum_{i=1}^{m} (f(x^{(i)}) - \ell(y, y^{(i)}))^2 + \lambda \|f\|_{\mathcal{H}}^2, \tag{2}$$

where $\|\cdot\|_{\mathcal{H}}$ is a norm over $\mathcal{H}$ and λ is a regularization parameter. We define $\mathcal{H}$ to be a reproducing kernel Hilbert space (RKHS) spanned by a real, symmetric positive definite kernel $k : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$.

[1] To simplify exposition, we omit the intercept here and discuss how it can be accounted for in the Appendix.

This enables us to learn from a rich class of functions and reduces the problem to kernel ridge regression (KRR) [4], leading to the following closed form solution.

Proposition 2.1. *For every $y \in \mathcal{Y}$, an optimal solution for the regularized least squares problem can be expressed as*

$$\hat{R}(y|\cdot) = \sum_{i=1}^{m} w_i(\cdot)\ell(y, y^{(i)}), \tag{3}$$

where $w_i(x)$ is the i-th element of $w(x) := (K + \lambda I)^{-1}v(x)$. Here $K := [k(x^{(i)}, x^{(j)})]_{ij}, \forall i, j \in \{1, \ldots, m\}$ is the $m \times m$ gram matrix and $v(x) := [k(x, x^{(i)})]_{i=1}^{m}$ is a $m \times 1$ vector.

Because the above holds true for all $y \in \mathcal{Y}$, we have a complete characterization of $\hat{R}(\cdot|x)$ as the weighted sum of the individual loss functions $\{\ell(\cdot, y^{(i)})\}_{i=1}^{m}$ induced by $y^{(1)}, \ldots, y^{(m)}$, with weights $w(x)$ that only depend on x and $x^{(1)}, \ldots, x^{(m)}$. This property arises from the fact that KRR is a *linear smoother* [5]. Such linear combinations of loss functions also arise in instance-based learning methods like kNN. However, these methods typically rely on a weighting of neighboring samples, whereas (3) is derived from least squares estimation over the whole data set. Using KRR provides added modeling flexibility through kernels, which can either be chosen based on domain knowledge or learned from data [6].

For comparison with existing structured prediction methods, it is useful to think of $\hat{R}(y|x)$ as an alternative scoring function to that of M^{3}N and SSVM, which also admits a linear form $\langle w, \phi(x, y)\rangle$. The key difference is that we do not define a joint feature map $\phi(x, y)$, but instead have a linear expression in which x and y are decoupled. This avoids the pre-image problem of finding a y that maximizes $\langle w, \phi(x, y)\rangle$ when predicting an output, which is nontrivial in general [7]. Another related method is Kernel Dependency Estimation (KDE) [8] and its extensions [9], which are also based on KRR. However, KDE differs from our approach in that regression is used to learn a mapping from inputs to a feature space associated with $\mathcal{Y}$, whose output is then mapped from the feature space back to $\mathcal{Y}$ by solving a pre-image problem. In place of the pre-image problem, our method requires solving an auxiliary optimization problem, which we discuss in the next section.

2.3 Prediction

Given an input x, we first compute $w(x) := (K + \lambda I)^{-1}v(x)$. This is done by first forming $v(x) \in \mathbb{R}^m$ and then either solving a linear system, or multiplying $v(x)$ with the inverted matrix if already pre-computed in training. Both can be done efficiently in moderately sized problems.[2] A prediction $y^* = \hat{h}(x)$ is computed by solving an auxiliary optimization problem that minimizes the estimated conditional risk $\hat{R}(y|x)$,

$$y^* \in \arg\min_{y \in \mathcal{Y}} \sum_{i=1}^{m} w_i(x)\ell(y, y^{(i)}). \tag{4}$$

The difficulty of this problem crucially depends on the geometry of $\mathcal{Y}$, the choice of ℓ and in some cases the signs of $w(x)$. In Section 3, we will apply LS-ECRM to a range of settings and provide conditions under which this auxiliary problem can solved efficiently.

3 Applications

3.1 Special case: $\mathcal{Y} = \{-1, 1\}$

Before we consider more complex settings, it is useful to understand what it means to apply LS-ECRM to the simple case of binary classification. With the zero-one loss, $\ell(y, y') = \mathbb{1}(y \neq y')$, $R(y|x)$ is simply the conditional probability of misclassification, $\mathbb{P}(y \neq Y|x)$. Minimizing its least squares estimate, we obtain the following decision rule from (4),

$$h(x) = 1 \text{ iff } \sum_{i=1}^{m} (\mathbb{1}(y^{(i)} = 1) - \mathbb{1}(y^{(i)} = -1))w_i(x) \geq 0. \tag{5}$$

[2] In large-scale problems, low-rank approximation of K [10] or a divide-and-conquer approach for KRR [11] can be considered.

We now show that this is equivalent to Regularized Least Squares Classification (RLSC)[3] [12]. In RLSC, one posits that the classification rule is of the form $h(x) = \text{sgn}(f(x))$ and estimates an f^* by solving a regularized least squares problem: $f^*(\cdot) \in \arg\min_{f \in \mathcal{H}} \sum_{i=1}^m (f(x^{(i)}) - y^{(i)})^2 + \lambda \|f\|_{\mathcal{H}}^2$. Here $\mathcal{H}$ is a RKHS associated with a kernel k. Despite using the square loss in place of the more common hinge loss of SVM, RLSC has been shown to achieve similar performance as SVM in practice [14]. The following proposition states that $\text{sgn}(f^*(x))$ is equivalent to (5).

Proposition 3.1. *Let $w(x)$ be defined as in Proposition 2.1. Then, sgn$(f^*(x)) = 1$ if and only if $\sum_{i=1}^m (\mathbb{1}(y^{(i)} = 1) - \mathbb{1}(y^{(i)} = -1)) w_i(x) \geq 0$.*

Thus LS-ECRM can be viewed as a generalization of RLSC. From the equivalence above, we can also interpret RLSC as minimizing the least squares estimated misclassification probability.

3.2 Predicting binary-valued vectors

In this section, we consider $\mathcal{Y} \subseteq \{0,1\}^d$ and any loss function $\ell(y, y')$ that can be written as the sum of functions of pairwise elements in y and y'. Specifically, we define $\mathcal{Y}$ to be a set of points in $\{0,1\}^d$ that satisfy the following set of linear constraints,

$$\begin{aligned} a_{i_1}^T y &\leq b_{i_1}, \forall i_1 \in I_1, \\ a_{i_2}^T y &\geq b_{i_2}, \forall i_2 \in I_2, \\ a_{i_3}^T y &= b_{i_3}, \forall i_3 \in I_3, \end{aligned} \tag{6}$$

where I_1, I_2, I_3 are (possibly empty) disjoint sets of indices such that their union is $\{1, \ldots, n\}$, and $a_i \in \mathbb{R}^d, b_i \in \mathbb{R}, \forall i$. In other words, there are a total of n linear constraints, each can either be an inequality or an equality. This characterization of $\mathcal{Y}$ is fairly general as many objects of interest, including matchings, permutations and label assignments in graphs, can be represented as such. For convenience, we will define $A := [a_1 \ldots a_n]^T$ to be the *constraint matrix* and $b := [b_1 \ldots b_n]^T$. One property of matrices that will be useful to us later is *total unimodularity*, defined as follows.

Definition (Total Unimodularity) A matrix A is totally unimodular if every square submatrix of A has determinant $0, -1$ or 1.

Now consider additive pairwise losses of the form, $\ell(y, y') = \sum_{j=1}^d \sum_{k=1}^d \ell_{jk}(y_j, y'_k)$. For example, the Hamming loss corresponds to $\ell_{jk}(y_j, y'_k) := \frac{1}{d} \mathbb{1}(j = k) \mathbb{1}(y_j \neq y'_k)$. Under this loss and our definition of $\mathcal{Y}$, the auxiliary problem in (4) is a discrete optimization problem,

$$\begin{aligned} \text{minimize} \quad & \sum_{i=1}^m \left(\sum_{j=1}^d \sum_{k=1}^d \ell_{jk}(y_j, y_k^{(i)}) \right) w_i(x) \\ \text{subject to} \quad & y \in \{0,1\}^d \cap \mathcal{Z} \end{aligned} \tag{7}$$

where $\mathcal{Z}$ is defined to be a polyhedron in $\mathbb{R}^d$ that is characterized by all the linear constraints in (6). This problem is $\mathcal{NP}$-complete in general. In the following theorem, we provide sufficient conditions under which the problem can be solved efficiently by exact linear programming relaxation.

Theorem 3.2. *If A is totally unimodular and $b \in \mathbb{Z}^n$, then for any $w(x) \in \mathbb{R}^m$, an optimal solution of the auxiliary problem can be found by solving the following linear program,*

$$\begin{aligned} \text{minimize} \quad & \sum_{j=1}^d \left(\sum_{k=1}^d \sum_{i=1}^n (\ell_{jk}(1, y_k^{(i)}) - \ell_{jk}(0, y_k^{(i)})) w_i(x) \right) y_j \\ \text{subject to} \quad & 0 \leq y_j \leq 1, \forall j = 1, \ldots, d \\ & y \in \mathcal{Z} \end{aligned} \tag{8}$$

Total unimodularity is an important property in combinatorial optimization. There are many results on classes of matrix that satisfy this property (see, e.g., [15]). In what follows, we will specialize this method to two problems and show how Theorem 3.2 can be applied, leading to new algorithms.

[3] RLSC is closely related to LS-SVM [13], which differs only in that it includes an unpenalized intercept.

3.2.1 Example 1: Hierarchical Multilabel Classification (HMC)

In HMC, we have a set of labels $\mathcal{V} = \{1, \ldots, d\}$ organized in a hierarchical structure (e.g., a hierarchy of topics in text classification). The goal is to predict a subset of $\mathcal{V}$ that corresponds to an input x. Let $y_j \in \{0, 1\}, j = 1, \ldots, d$ denote whether each label y_j is chosen (1) or not (0). In addition to choosing a subset of $\mathcal{V}$, we consider a variant of this problem where y must also satisfy the following hierarchical constraints:

- For each $j \in \mathcal{V}$: if $y_j = 1$, then $y_k = 1, \forall k \in \mathcal{P}(j)$. Here $\mathcal{P}(j)$ denotes the set of immediate parent labels under which j belongs in the hierarchy.

The hierarchy is commonly represented as a tree or (more generally) a directed acyclic graph (DAG), with each node being a label in $\mathcal{V}$ and each arc (k, j) encoding a parent-child relation, $k \in \mathcal{P}(j)$. Formally, we define the DAG as $G = (\mathcal{V}, \mathcal{A})$, where $\mathcal{V}$ is defined as above and $(k, j) \in \mathcal{A}$ iff $k \in \mathcal{P}(j)$. The resulting output space $\mathcal{Y}$ can be succintly described with $|\mathcal{A}|$ linear constraints (in addition to $y \in \{0, 1\}^d$),

$$y_j \leq y_k, \forall (k, j) \in \mathcal{A}.$$

A common choice of loss function in HMC is the Hamming loss, $\ell(y, y') = \frac{1}{d} \sum_{j=1}^{m} \mathbb{1}(y_j \neq y'_j)$, which is simply the fraction of incorrect labels. This is a special case of the additive pairwise losses defined in the preceding section. We show that the auxiliary problem can be solved efficiently below.

Proposition 3.3. *For any DAG, $G = (\mathcal{V}, \mathcal{A})$, the auxiliary problem for HMC with Hamming loss is equivalent to the following linear program,*

$$\begin{aligned} \text{minimize} \quad & \sum_{j=1}^{d} \left(\sum_{i=1}^{m} (\mathbb{1}(y_j^{(i)} = 0) - \mathbb{1}(y_j^{(i)} = 1)) w_i(x) \right) y_j \\ \text{subject to} \quad & y_j \leq y_k, \forall (k, j) \in \mathcal{A} \\ & 0 \leq y_j \leq 1, \forall j = 1, \ldots, d \end{aligned} \tag{9}$$

Given $w(x)$, the cost coefficients of the linear program can be computed in $O(md)$. The size of the linear program (in terms of variable and constraint counts) scales linearly with d and $|\mathcal{A}|$, but is otherwise independent of m. In practice, linear programs can be solved very efficiently with state-of-the-art optimization solvers [16]. As we will see in Section 6, this makes the method scale well to large hierarchies.

3.2.2 Example 2: Multilabel Ranking

Suppose we are interested in predicting a ranking over all labels, rather than only choosing a subset. We consider a setting where the training set consists of complete permutations $\sigma^{(1)}, \ldots, \sigma^{(m)}$ over the label set $\mathcal{V}$ and their associated inputs. The goal is to learn to predict a permutation σ given an input x, where $\sigma(j)$ indicates the rank of label j for every $j = 1, \ldots, d$. To measure the loss, we will use the normalized *Spearman's footrule distance*, $\tilde{\ell}(\sigma, \sigma') = \frac{1}{2d(d+1)} \sum_{j=1}^{d} |\sigma(j) - \sigma'(j)|$, which sums the absolute differences of the two ranks over all labels and normalizes it to within $[0, 1]$. By representing σ as a binary vector, we will show that the auxiliary problem can be solved by exact linear programming relaxation.

Let us define $y_{j,k} \in \{0, 1\}, \forall j, k \in \{1, \ldots, d\}$, as a vector in $\mathbb{R}^{d^2}$ that corresponds to some σ as follows: $(\forall j, k)\ y_{jk} = 1$ iff $\sigma(j) = k$. The set of such vectors (each corresponding to a distinct σ), denoted $\mathcal{Y} \subseteq \{0, 1\}^{d^2}$, are exactly characterized by the following linear constraints.

$$\begin{aligned} &\textstyle\sum_{k=1}^{d} y_{jk} = 1, \forall j = 1, \ldots, d \\ &\textstyle\sum_{j=1}^{d} y_{jk} = 1, \forall k = 1, \ldots, d \end{aligned} \tag{10}$$

This can be interpreted as the set of perfect matchings in a complete bipartite graph. With this representation, we can equivalently write $\tilde{\ell}$ as $\ell(y, y') = \frac{1}{2d(d+1)} \sum_{j,k,l} |k - l| \mathbb{1}(y_{jk} = 1, y'_{jl} = 1)$, where each summation over j, k, l is from 1 to d. This is again a special case of the additive pairwise losses defined earlier. Together with the fact that the constraint matrix associated with bipartite matchings is known to be totally unimodular [15], we can apply Theorem 3.2 to reduce the auxiliary problem to the following min-cost *assignment problem*.

Proposition 3.4. *The auxiliary problem for multilabel ranking with Spearman's footrule distance is equivalent to the following linear program,*

$$\begin{aligned} \text{minimize} \quad & \sum_{j=1}^{d}\sum_{k=1}^{d}\left(\sum_{l=1}^{d}|k-l|\sum_{i=1}^{m}\mathbb{1}(\sigma^{(i)}(j)=l)w_i(x)\right)y_{jk} \\ \text{subject to} \quad & \textstyle\sum_{k=1}^{d} y_{jk} = 1, \forall j = 1,\ldots,d \\ & \textstyle\sum_{j=1}^{d} y_{jk} = 1, \forall k = 1,\ldots,d \\ & y_{jk} \geq 0, \forall j = 1,\ldots,d, \forall k = 1,\ldots,d \end{aligned} \tag{11}$$

4 Extensions

4.1 Predicting continuous-valued vectors

We briefly describe how LS-ECRM can be extended to predict a vector in a convex set $\mathcal{Y}$, under a loss function $\ell(y, y')$ that is convex in y (e.g., L_1 and L_∞ losses). In principle, Proposition 2.1 still applies but the auxiliary problem in (4) is now a continuous optimization problem. Also, the convexity of $\ell(\cdot, y')$ implies that $R(\cdot|x)$ is convex for every x. However, this need not be true for $\hat{R}(y|x) = \sum_{i=1}^{m} w_i(x)\ell(y, y^{(i)})$ if some element of $w(x)$ is negative, which can result in a non-convex auxiliary problem. One approach is to solve the problem approximately by minimizing its convex upper bound $\hat{R}^+(y|x) := \sum_{i=1}^{m} \max\{w_i(x), 0)\}\ell(y, y^{(i)})$, which implies that

$$\min_{y\in\mathcal{Y}} \hat{R}(y|x) \leq \min_{y\in\mathcal{Y}} \hat{R}^+(y|x). \tag{12}$$

While we do not provide examples of application in this paper, we note that the problem of learning a mapping from inputs to a convex set has been studied in the more general setting of conditional stochastic optimization using local estimators [17][18] such as kNN and kernel smoothing.

4.2 Learning additive models of conditional risk

As seen in examples of the previous section, many loss functions considered in structured prediction are additive over substructures of $\mathcal{Y}$. It is often desirable to decompose the learning problem over these substructures, so that any local features can be exploited for better generalization performance. Here we show how LS-ECRM can be extended for this purpose. Consider $\ell(y, y') = \sum_{S\in\mathcal{S}} \ell_S(y_S, y'_S)$, where each ℓ_S is defined over a subset of elements in y, y' that correspond to the index set $S \subseteq \{1,\ldots,d\}$, i.e., $y_S := (y_j : j \in S)$ (likewise, for y'_S). As a result, we can also express the conditional risk function in additive form, $R(y|x) = \sum_{s\in\mathcal{S}} R_S(y_S|x)$ with $R_S(y_S|x) := \mathbb{E}[\ell_S(y_S, Y)|x]$, and then estimate $R_S(y_S|x)$ separately for every S.

This decomposition is useful in several scenarios. For example, with respect to each S, we may: (i) tune the regularization and kernel parameters in (2) (ii) apply feature selection so that $\hat{R}_S(y_S|x)$ depends only on a subset of inputs most relevant for predicting y_S. By solving the regularized least squares problem in (2), we obtain for every S, $\hat{R}_S(y_S|x) = \sum_{i=1}^{m} w_{S,i}(x)\ell(y_S, y_S^{(i)})$, where $w_S(x)$ is a weight vector specific to S. These results can then be combined to form

$$\hat{R}(y|x) = \sum_{S\in\mathcal{S}} \hat{R}_S(y_S|x) = \sum_{i=1}^{m}\sum_{S\in\mathcal{S}} w_{S,i}(x)\ell(y_S, y_S^{(i)}).$$

and the auxiliary problem can be solved similarly as in (4). Theorem 3.2, Proposition 3.3 and 3.4 can also be extended to hold in this case if the losses are in the form described in Section 3.2. On the other hand, such decompositions incur additional computational costs because a separate weight vector has to be computed for each S. Deciding the granularity of S is thus a matter of tradeoff between predictive and computational performance, and largely depends on the application.

5 Statistical properties of LS-ECRM

We analyze the statistical properties of LS-ECRM under the assumption that $R(y|\cdot) \in \mathcal{H}$ for a parametric class of functions $\mathcal{H}$. Specifically, for every $y \in \mathcal{Y}$, we assume that there exists some

$\beta(y) \in \mathbb{R}^q$ such that $R(y|x) = \beta(y)^T \phi(x)$, where $\phi : \mathcal{X} \mapsto \mathbb{R}^q$ is a mapping from the input space to a Euclidean *feature space* with finite dimension q. Thus (2) can be reduced to linear regression in the feature space,[4]

$$\hat{\beta}(y) \in \arg\min_{\beta \in \mathbb{R}^q} \sum_{i=1}^{m} \left(\beta^T \phi(x^{(i)}) - \ell(y, y^{(i)}) \right)^2 + \lambda \|\beta\|_2^2, \tag{13}$$

We will only consider the setting where $\lambda = 0$ and the *design matrix*, $\Phi := [\phi(x^{(1)}) \ \cdots \ \phi(x^{(m)})]^T \in \mathbb{R}^{m \times q}$ is observed and non-random, which is analogous to the *fixed design* setting in linear regression.[5] Under these assumptions, we derive the following excess risk bound.

Theorem 5.1. *Let $\hat{h}$ and h^* be a LS-ECRM and Bayes optimal predictor, respectively. Suppose that (i) $\mathcal{Y}$ is finite (ii) $\sup_{y,y' \in \mathcal{Y}} \ell(y, y') \leq L$ (iii) $\mathbb{E}[\|\phi(X)\|_2] \leq B$. Then*

$$\mathbb{E}[\ell(\hat{h}(X), Y)] - \mathbb{E}[\ell(h^*(X), Y)] \leq 4BL \sqrt{\frac{r(4 \log |\mathcal{Y}| + 3)}{\nu m}}, \tag{14}$$

where $r \leq q$ and ν are the rank and the smallest eigenvalue of $\frac{1}{m}\Phi^T\Phi$, respectively.

For example, in the HMC problem with Hamming loss, we have $L = 1$ and $|\mathcal{Y}| \leq 2^d$. This gives a bound of $O(\sqrt{d/m})$, where d is the number of labels.

6 Experiments

6.1 Setup

We apply LS-ECRM in HMC using the formulation in Section 3.2.1 and compare it with several existing algorithms, each based on a different paradigm: (i) BR-SVM: A *binary relevance* approach: a SVM classifier is trained for each node. The predictions are imputed from the bottom up to satisfy the hierarchical constraints. (ii) H-SVM: A hierarchical approach [20]: at each node, a SVM classifier is trained using only samples in which its parents are positively labeled. Node labels are predicted recursively from the top down, terminating at any node that is negatively labeled. (iii) Clus-HMC [21]: A decision tree approach.

Each method is evaluated on HMC benchmark data sets from three domains: text (enron, rcv1v2), image (imclef07a, imclef07d) and functional genomics (pheno_go, spo_go, gasch1_go, gasch2_go). A summary of these data sets are available in the Appendix. For consistency, we do not apply any feature selection. Preprocessing is only done for the genomics data sets, where we trim the original hierarchy with over 4000 nodes down to a few hundreds. We implement our method in Julia [22] and solve the auxiliary linear programs with Clp [23]. Both BR-SVM and H-SVM are implemented using LibSVM [24]. For Clus-HMC, bagging is used to train an ensemble of decision trees, which we have observed offer substantial accuracy gains over a single decision tree. LS-ECRM, BR-SVM and H-SVM all use the RBF kernel. The parameters are tuned on the training sets by grid search with cross validation. All tests are run on a quad-core 3.5GHz CPU with 8GB RAM.

6.2 Results

For each data set, we compute the AUPRCs (area under precision-recall curve) for all methods within the intersection of their operating regions and the highest F-scores in their respective PR curves. This is shown in Table 1. To find out how the methods scale computationally with hierarchy size, we evaluate the training and prediction time on gasch1_go over label counts ranging from 10 to 512, obtained by truncating the original hierarchy. The results are presented in Figure 1.

As Table 1 shows, LS-ECRM and BR-SVM perform the best overall in terms of both metrics, leading the other two methods by significant margins in several data sets. While LS-ECRM has only a slight lead over BR-SVM in predictive performance, its training and and prediction time can be several orders of magnitude faster, as Figure 1 shows. Whereas BR-SVM requires training a separate classifier at each node, training LS-ECRM only involves forming the kernel matrix and computing

[4] The intercept can be included with a constant, $\phi(x) = [\bar{\phi}(x) \ 1]$, which is unpenalized if $\lambda = 0$.

[5] This greatly simplifies the analysis, though it can be extended to more general settings [19].

Table 1: The F-scores and AUPRCs of each method on various data sets. For each data set, the rankings of methods by each metric are shown in parentheses, and the best results are in bold.

	LS-ECRM		BR-SVM		H-SVM		Clus-HMC	
Data set	F-score	AUPRC	F-score	AUPRC	F-score	AUPRC	F-score	AUPRC
enron	**0.696** (1)	**0.546** (1)	0.692 (2)	0.543 (2)	0.671 (3)	0.517 (3)	0.636 (4)	0.474 (4)
rcv1v2	0.758 (2)	0.534 (2)	**0.772** (1)	**0.549** (1)	0.743 (3)	0.516 (3)	0.656 (4)	0.426 (4)
imclef07a	0.324 (2)	0.102 (3)	**0.335** (1)	0.104 (2)	0.319 (4)	0.094 (4)	0.321 (3)	**0.108** (1)
imclef07d	0.444 (2)	0.121 (2)	0.431 (3)	0.117 (3)	0.423 (4)	0.113 (4)	**0.470** (1)	**0.136** (1)
pheno_go	**0.412** (1)	**0.210** (1)	0.409 (2)	0.208 (2)	0.397 (4)	0.145 (4)	0.409 (2)	0.203 (3)
spo_go	0.392 (2)	0.198 (2)	0.392 (2)	0.198 (2)	0.382 (4)	0.140 (4)	**0.402** (1)	**0.209** (1)
gasch1_go	0.435 (2)	0.249 (2)	**0.444** (1)	**0.261** (1)	0.399 (4)	0.165 (4)	0.421 (3)	0.233 (3)
gasch2_go	**0.431** (1)	**0.237** (1)	0.427 (2)	**0.237** (1)	0.367 (4)	0.137 (4)	0.416 (3)	0.226 (3)
Avg rank	1.625	1.750	1.750	1.750	3.750	3.750	2.625	2.500

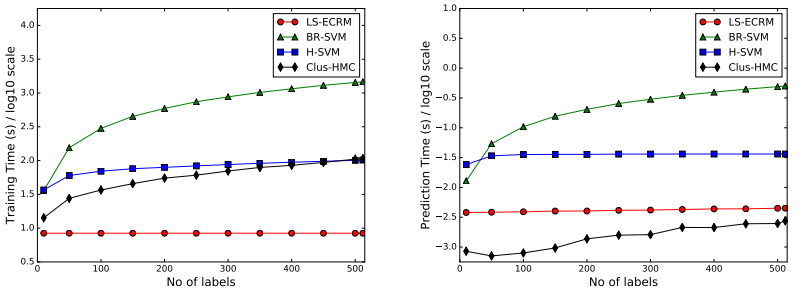

(a) Total training time

(b) Average prediction time per instance

Figure 1: Training and prediction time ($\log_{10}$ scale) for the gasch1_go dataset under varying hierarchy sizes. The training and test sets have 2477 and 1281 instances, respectively.

an inversion, which is independent of the hierarchy size. This is a unique property not shared by other methods under our evaluation. Predicting with LS-ECRM requires solving linear programs, which we have observed empirically tend to terminate in few iterations without exhaustive search, in spite of growing hierarchy size. This gives it a clear advantage over BR-SVM, which does not take into account of the problem structure to search efficiently. H-SVM also has a flattening trend in both training and prediction time, which can be explained by its top-down hierarchical approach: even as the hierarchy grows, the number of positively labeled instances generally diminish (for a fixed sample size) as one descends down the hierarchy, resulting in early termination of the algorithm. In terms of prediction time, Clus-HMC has a clear lead over the rest, though its gap with LS-ECRM narrows with increasing label counts. Its predictive performance, however, trails behind in several data sets.

7 Conclusions

We have developed a new approach to structured prediction based on least squares estimated conditional risk minimization. Our approach treats the problem as one of learning a conditional risk function, which is then minimized by solving an auxiliary problem to predict an output. We showed how kernel methods can be used to learn from a rich class of hypotheses without introducing a joint feature map, thus avoiding the pre-image problem of existing methods. By using exact linear programming relaxations, we demonstrated that the auxiliary problem can be solved efficiently in several important cases, including hierarchical multilabel classification (HMC) and multilabel ranking with additive pairwise losses. Empirical evaluations on benchmark data sets showed that the proposed HMC algorithm compares favorably with existing methods and scales well computationally to large hierarchies, most notably having a training time that is independent of the hierarchy size. The generality of our conditional risk minimization framework and its compatibility with a variety of loss functions makes the approach potentially applicable to many problems.

References

[1] Lafferty, J., McCallum, A., & Pereira, F. C. (2001). Conditional random fields: Probabilistic models for segmenting and labeling sequence data.

[2] Taskar, B., Guestrin, C., & Koller, D. (2003). Max-margin Markov networks. Advances in Neural Information Processing Systems 16, 25-32.

[3] Tsochantaridis, I., Joachims, T., Hofmann, T., & Altun, Y. (2005). Large margin methods for structured and interdependent output variables. Journal of Machine Learning Research (JMLR), 6, 1453-1484.

[4] Saunders, C., Gammerman, A., & Vovk, V. (1998). Ridge regression learning algorithm in dual variables. Proceedings of the 15th International Conference on Machine Learning (pp. 515-521). Morgan Kaufmann.

[5] Buja, A., Hastie, T., & Tibshirani, R. (1989). Linear smoothers and additive models. The Annals of Statistics, 453-510.

[6] Cortes, C., Mohri, M., & Rostamizadeh, a. (2009). L2 regularization for learning kernels. Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (pp. 109-116). AUAI Press.

[7] Gärtner, T., & Vembu, S. (2009). On structured output training: hard cases and an efficient alternative. Machine Learning, 76(2-3), 227-242.

[8] Weston, J., Chapelle, O., Elisseeff, A., Sch, B., & Vapnik, V. (2002). Kernel dependency estimation. Advances in Neural Information Processing Systems (NIPS), 873 - 880. Retrieved from http://research.yahoo.com/pub/1178

[9] Cortes, C., Mohri, M., & Weston, J. (2005). A general regression technique for learning transductions. Proceedings of the 22nd International Conference on Machine Learning, 153-160.

[10] Kumar, S., Mohri, M., & Talwalkar, A. (2009). Ensemble nystrom method. Advances in Neural Information Processing Systems (pp. 1060-1068).

[11] Zhang, Y., Duchi, J., & Wainwright, M. (2013). Divide and conquer kernel ridge regression. Conference on Learning Theory (pp. 592-617).

[12] Rifkin, R., Yeo, G., & Poggio, T. (2003). Regularized least-squares classification. Advances in Learning Theory: Methods, Models and Applications, Nato Science Series Sub Series III Computer and Systems Sciences, 190, 131-154.

[13] Suykens, J. A., & Vandewalle, J. (1999). Least squares support vector machine classifiers. Neural processing letters, 9(3), 293-300.

[14] Gestel, T. van, Suykens, J. A. K., Baesens, B., Viaene, S., Vanthienen, J., Dedene, G., Moor, B. de, et al. (2004). Benchmarking Least Squares Support Vector Machine Classifiers. Machine Learning, 54(1), 5-32.

[15] Schrijver, A. (1998). Theory of linear and integer programming. John Wiley & Sons.

[16] Mittelmann, H. (2015). Benchmarks for optimization software. Retrieved from `http://plato.asu.edu/bench.html`.

[17] Hannah, L., Powell, W., & Blei, D. M. (2010). Nonparametric density estimation for stochastic optimization with an observable state variable. Advances in Neural Information Processing Systems (pp. 820-828).

[18] Bertsimas, D., & Kallus, N. (2015). From Predictive to prescriptive analytics. arXiv preprint arXiv:1402.5481.

[19] Hsu, D., Kakade, S. M., & Zhang, T. (2014). Random design analysis of ridge regression. Foundations of Computational Mathematics, 14(3), 569-600.

[20] Cesa-Bianchi, N., Gentile, C., & Zaniboni, L. (2006). Incremental algorithms for hierarchical classification. The Journal of Machine Learning Research, 7, 31-54.

[21] Vens, C., Struyf, J., Schietgat, L., Džeroski, S., & Blockeel, H. (2008). Decision trees for hierarchical multi-label classification. Machine Learning, 73(2), 185-214.

[22] Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2014). Julia: A fresh approach to numerical computing. arXiv preprint arXiv:1411.1607.

[23] COIN-OR Linear Programming Solver. (2016). Retrieved from `https://projects.coin-or.org/Clp`

[24] Chang, C. C., & Lin, C. J. (2011). LIBSVM: a library for support vector machines. ACM Transactions on Intelligent Systems and Technology (TIST), 2(3), 27.

1 Appendix A: Proofs

1.1 Proof of Proposition 2.1

Proposition 2.1. *For every $y \in \mathcal{Y}$, an optimal solution for the regularized least squares problem can be expressed as*

$$\hat{R}(y|\cdot) = \sum_{i=1}^{m} w_i(\cdot)\ell(y, y^{(i)}), \tag{1}$$

where $w_i(x)$ is the i-th element of $w(x) := (K+\lambda I)^{-1}v(x)$. Here $K := [k(x^{(i)}, x^{(j)})]_{ij}, \forall i, j \in \{1, \ldots, m\}$ is the $m \times m$ gram matrix and $v(x) := [k(x, x^{(i)})]_{i=1}^{m}$ is a $m \times 1$ vector.

Proof. Recall that the optimization problem is the following.

$$\hat{R}(y|\cdot) \in \arg\min_{f \in \mathcal{H}} \sum_{i=1}^{m} (f(x^{(i)}) - \ell(y, y^{(i)}))^2 + \lambda \|f\|_{\mathcal{H}}^2,$$

For any fixed y, the *representer's theorem* [1] states that any optimal solution of the problem lies in the span of $\{k(\cdot, x^{(i)}) : i = 1, \ldots, m\}$, i.e., there exists some $\alpha \in \mathbb{R}^m$ such that $\hat{R}(y|\cdot) = \sum_{i=1}^{m} \alpha_i k(\cdot, x^{(i)})$. This also implies that $\|\hat{R}(y|\cdot)\|_{\mathcal{H}}^2 = \alpha^T K \alpha$, where K is defined above. Thus it suffices to consider an equivalent optimization problem over α,

$$\alpha_y^* \in \arg\min_{\alpha \in \mathbb{R}^m} \sum_{i=1}^{m} \left(\sum_{j=1}^{m} \alpha_j k(x^{(i)}, x^{(j)}) - \ell(y, y^{(i)}) \right)^2 + \lambda \alpha^T K \alpha. \tag{2}$$

Denoting $L_y := [\ell(y, y^{(1)}) \ \cdots \ \ell(y, y^{(m)})]$ as the vector of observed losses, this can be written compactly as

$$\alpha_y^* \in \arg\min_{\alpha \in \mathbb{R}^m} \|K\alpha - L_y\|_2^2 + \lambda \alpha^T K \alpha. \tag{3}$$

Since the problem is convex (which follows from k being a symmetric, positive definite kernel), the first order optimality conditions are necessary and sufficient. Taking the derivative of the objective function and equating it to zero, we have

$$2K^T K \alpha_y^* - 2KL_y + 2\lambda K \alpha_y^* = 0$$
$$\alpha_y^* = (K + \lambda I)^{-1} L_y.$$

For any $x \in \mathcal{X}$, we know that $\hat{R}(y|x) = \sum_{i=1}^{m} \alpha_{y,i}^* k(x, x^{(i)}) = \alpha_y^{*T} v(x)$. Substituting the above into this expression, we obtain $L_y^T (K + \lambda I)^{-1} v(x) = \sum_{i=1}^{m} w_i(x) \ell(y, y^{(i)})$ as desired. □

1.2 Proof of Proposition 3.1

Proposition 3.1. *Let $w(x)$ be defined as in Proposition 2.1. Then, sgn$(f^*(x)) = 1$ if and only if $\sum_{i=1}^{m} (\mathbb{1}(y^{(i)} = 1) - \mathbb{1}(y^{(i)} = -1)) w_i(x) \geq 0$.*

Proof. We can solve for f^* by applying the representer's theorem as in the proof of Proposition 2.1,

$$f^*(x) = \sum_{i=1}^{m} w_i(x) y^{(i)} = \sum_{i=1}^{m} w_i(x) (\mathbb{1}(y^{(i)} = 1) - \mathbb{1}(y^{(i)} = -1)).$$

Since sgn$(f^*(x)) = 1$ iff $f^*(x) \geq 0$, the claim is proven. □

1.3 Proof of Theorem 3.2

Theorem 3.2. *If A is totally unimodular and $b \in \mathbb{Z}^n$, then for any $w(x) \in \mathbb{R}^m$, an optimal solution of the auxiliary problem can be found by solving the following linear program,*

$$\begin{aligned} \text{minimize} \quad & \sum_{j=1}^{d} \left(\sum_{k=1}^{d} \sum_{i=1}^{n} (\ell_{jk}(1, y_k^{(i)}) - \ell_{jk}(0, y_k^{(i)})) w_i(x) \right) y_j \\ \text{subject to} \quad & 0 \leq y_j \leq 1, \forall j = 1, \ldots, d \\ & y \in \mathcal{Z} \end{aligned} \tag{4}$$

Proof. First, we show that the auxiliary problem is equivalent to an integer linear program. For that purpose, it is useful to express the loss function algebraically in terms of y: because $y_j \in \{0,1\}$, we have $\ell_{jk}(y_j, y'_k) = \ell_{jk}(1, y'_k)y_j + \ell_{jk}(0, y'_k)(1 - y_j), \forall j, \forall k$. As a shorthand notation, let us denote $\bar{\ell}^{(i)}_{jk} := \ell_{jk}(1, y^{(i)}_k)$ and $\dot{\ell}^{(i)}_{jk} := \ell_{jk}(0, y^{(i)}_k)$. Then,

$$\begin{aligned}
\min_{y\in\mathcal{Y}} \sum_{i=1}^{m} \left(\sum_{j=1}^{d}\sum_{k=1}^{d} \ell_{jk}(y_j, y^{(i)}_k) \right) w_i(x) &= \min_{y\in\mathcal{Y}} \sum_{i=1}^{m} \left(\sum_{j=1}^{d}\sum_{k=1}^{d} \bar{\ell}^{(i)}_{jk} y_j + \dot{\ell}^{(i)}_{jk}(1 - y_j) \right) w_i(x) \\
&= \min_{y\in\mathcal{Y}} \sum_{j=1}^{d} \left(\sum_{k=1}^{d}\sum_{i=1}^{m} (\bar{\ell}^{(i)}_{jk} - \dot{\ell}^{(i)}_{jk}) y_j w_i(x) \right) + \sum_{j=1}^{d}\sum_{k=1}^{d}\sum_{i=1}^{m} \dot{\ell}^{(i)}_{jk} w_i(x) \\
&= C + \min_{y\in\mathcal{Y}} \sum_{j=1}^{d} \left(\sum_{k=1}^{d}\sum_{i=1}^{m} (\bar{\ell}^{(i)}_{jk} - \dot{\ell}^{(i)}_{jk}) w_i(x) \right) y_j.
\end{aligned}$$

Here $C := \sum_{j,k,i} \dot{\ell}^{(i)}_{jk} w_i(x)$ is a constant that does not depend on the decision variables. Thus the problem is equivalent to minimizing over a linear objective function subject to $y \in \{0,1\}^d \cap \mathcal{Z}$.

Next, we establish sufficient conditions for the existence of an exact linear programming relaxation of the above integer linear program. This allows us to replace all constraints $y_j \in \{0,1\}$ with $0 \le y_j \le 1, \forall j$, such that the resulting linear program (LP) is guaranteed to contain an optimal solution that is also optimal for the original problem. The key is to show that the set of linear constraints that characterize the LP,

$$a^T_{i_1} y \le b_{i_1}, \forall i_1 \in I_1, \tag{5}$$
$$a^T_{i_2} y \ge b_{i_2}, \forall i_2 \in I_2, \tag{6}$$
$$a^T_{i_3} y = b_{i_3}, \forall i_3 \in I_3, \tag{7}$$
$$y_j \ge 0, \forall j = 1, \ldots, d, \tag{8}$$
$$y_j \le 1, \forall j = 1, \ldots, d, \tag{9}$$

can be expressed in a single constraint matrix that is totally unimodular (TU). Let $A' := \left[A^T | I | I\right]^T$ be the $(n + 2d) \times d$ constraint matrix that characterizes the LHS of (5)-(9), where $I \in \mathbb{R}^{d\times d}$ is an identity matrix that corresponds to either (8) or (9). Likewise, let $b' := [b^T | 0^T | 1^T]^T$ be a $(n + 2d) \times 1$ vector that characterizes the RHS, where 0^T and 1^T are $1 \times d$ row vectors of zeros and ones, respectively. By the assumption of the theorem, we know that A is TU. Our goal now is to establish that A' is also TU.

Lemma 1.2. *If A is a TU matrix, then (i) A^T is TU (ii) $[A|I]$ is TU.*

Proof. Both of these results are well known. (i) follows directly from the fact that $\det(B^T) = \det(B)$ for any square matrix B. A proof of (ii) is available in [2]. □

By (i) and (ii) in Lemma 1.2, $A \Rightarrow A^T \Rightarrow [A^T|I] \Rightarrow [A^T|I|I] \Rightarrow [A^T|I|I]^T := A'$ is a chain of operations that preserve TU, thus proving the claim that A' is TU.

We now show that this implies every vertex of the polytope characterized by (5)-(9) is integral. At a vertex, there exist d constraints out of (5)-(9) that hold with equality (by default, (7) is included), such that the corresponding constraint vectors are linearly independent. Let $\bar{B}$ and $\bar{b}$ be submatrices of A' and b', respectively, that correspond to these d rows of constraints. Then the vertex is the unique solution of these equalities, $\bar{B}^{-1}\bar{b}$. Note that $\bar{b}$ is integral because b is integral by assumption. Also, by total unimodularity of A', $\det(\bar{B})$ is either $0, -1$ or 1. Therefore, by Cramer's rule, $\bar{B}^{-1}\bar{b}$ is integral.

Together with the well-known fact that any LP defined over a nonempty polytope contains an optimal solution that is a vertex [3], we conclude that the LP in (4) has an optimal integer solution that is a vertex. As a result of the relaxation, it must also be optimal for the integer linear program. □

1.4 Proof of Proposition 3.3

Proposition 3.3. *For any DAG, $G = (\mathcal{V}, \mathcal{A})$, the auxiliary problem for HMC with Hamming loss is equivalent to the following linear program,*

$$
\begin{aligned}
\text{minimize} \quad & \sum_{j=1}^{d} \left(\sum_{i=1}^{m} (\mathbb{1}(y_j^{(i)} = 0) - \mathbb{1}(y_j^{(i)} = 1)) w_i(x) \right) y_j \\
\text{subject to} \quad & y_j \leq y_k, \forall (k,j) \in \mathcal{A} \\
& 0 \leq y_j \leq 1, \forall j = 1, \ldots, d
\end{aligned}
\tag{10}
$$

Proof. In order to apply Theorem 3.2, we will show that the set of constraints, $y_j - y_k \leq 0, \forall (k,j) \in \mathcal{A}$ can be represented by a constraint matrix A and coefficient vector b, such that A is totally unimodular (TU) and b is integral. More precisely, we want to represent the constraints in (10) as

$$
\begin{aligned}
& Ay \leq b \\
& 0 \leq y_j \leq 1, \forall j = 1, \ldots, d
\end{aligned}
$$

where $A \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{V}|}$ and $b \in \mathbb{R}^{|\mathcal{A}|}$. Clearly, $b := 0$ is integral. We now show that A^T is exactly characterized by the following lemma.

Lemma 1.3. *Suppose $B \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{A}|}$ is a matrix with each element B_{ja} corresponding to the node-arc pair (j,a), $j \in \mathcal{V}, a \in \mathcal{A}$, such that*[6]

$$
B_{ja} = \begin{cases} -1, & \text{if } j \text{ is the head of } a, \\ 1, & \text{if } j \text{ is the tail of } a, \\ 0, & \text{otherwise.} \end{cases}
$$

Then, B is TU.

Proof. A proof is available in Chapter 4 of [4]. □

It is easy to check that $A^T = B$, where B is defined in the lemma above. Lemma 1.2 implies that B is also TU. Therefore, we satisfy all conditions of Theorem 3.2 and by substituting $\ell_{jk}(y_j, y'_k) := \mathbb{1}(j = k)\mathbb{1}(y_j \neq y'_k)$, our claim is proven. □

1.5 Proof of Proposition 3.4

Proposition 3.4. *The auxiliary problem for multilabel ranking with Spearman's footrule distance is equivalent to the following linear program,*

$$
\begin{aligned}
\text{minimize} \quad & \sum_{j=1}^{d} \sum_{k=1}^{d} \left(\sum_{l=1}^{d} |k-l| \sum_{i=1}^{m} \mathbb{1}(\sigma^{(i)}(j) = l) w_i(x) \right) y_{jk} \\
\text{subject to} \quad & \textstyle\sum_{k=1}^{d} y_{jk} = 1, \forall j = 1, \ldots, d \\
& \textstyle\sum_{j=1}^{d} y_{jk} = 1, \forall k = 1, \ldots, d \\
& y_{jk} \geq 0, \forall j = 1, \ldots, d, \forall k = 1, \ldots, d
\end{aligned}
\tag{11}
$$

Proof. By noting that $\mathbb{1}(y_{jk} = 1, y'_{jl} = 1) = \mathbb{1}(y'_{jl} = 1) y_{jk}$, we can write the objective function of the auxiliary problem as follows,

$$
\begin{aligned}
\sum_{i=1}^{m} w_i(x) \ell(y, y^{(i)}) &= \frac{1}{2d(d+1)} \sum_{i=1}^{m} w_i(x) \sum_{j=1}^{m} \sum_{k=1}^{m} \sum_{l=1}^{m} |k-l| \mathbb{1}(y_{jl}^{(i)} = 1) y_{jk} \\
&= \frac{1}{2d(d+1)} \sum_{j=1}^{m} \sum_{k=1}^{m} \left(\sum_{l=1}^{m} |k-l| \sum_{i=1}^{m} \mathbb{1}(y_{jl}^{(i)} = 1) w_i(x) \right) y_{jk}
\end{aligned}
$$

By definition, $y_{jl}^{(i)} = 1$ iff $\sigma^{(i)}(j) = l$, so we obtain the objective in the proposition above (disregarding the normalization factor). The constraint matrix associated with perfect bipartite matchings is known to be totally unimodular [4]. We can thus apply Theorem 3.2 and relax the constraints so that $y_{jk} \in \{0, 1\}$ becomes $0 \leq y_{jk} \leq 1, \forall j, k$. The inclusion of constraints $y_{jk} \leq 1$ does not change the problem in (11), so they are redundant and can be dropped. □

[6] For an arc $a = (k, j)$, k is the *head* and j is the *tail*.

1.6 Proof of Theorem 5.1

Theorem 5.1. *Let $\hat{h}$ and h^* be a LS-ECRM and Bayes optimal predictor, respectively. Suppose that (i) $\mathcal{Y}$ is finite (ii) $\sup_{y,y' \in \mathcal{Y}} \ell(y, y') \leq L$ (iii) $\mathbb{E}[\|\phi(X)\|_2] \leq B$. Then*

$$\mathbb{E}[\ell(\hat{h}(X), Y)] - \mathbb{E}[\ell(h^*(X), Y)] \leq 4BL\sqrt{\frac{r(4\log|\mathcal{Y}| + 3)}{\nu m}}, \tag{12}$$

where $r \leq q$ and ν are the rank and the smallest eigenvalue of $\frac{1}{m}\Phi^T\Phi$, respectively.

The proof extends the analysis of fixed design (Φ is deterministic) linear regression, with the key extension being that the results hold uniformly over $y \in \mathcal{Y}$. It can be broken down into 3 steps.

1. Bound the LHS of (12) with $2\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\hat{\beta}(y) - \beta(y)\|_2\right] \mathbb{E}[\|\phi(X)\|_2]$.
2. Show that $\sup_{y\in\mathcal{Y}} \|\hat{\beta}(y) - \beta(y)\|_2^2 \leq \sup_{y\in\mathcal{Y}} \frac{\|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2^2}{\nu m}$.
3. Finally, bound $\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2\right]$ with maximal inequalities.

1.6.1 Proof of Theorem 5.1: Step 1

Lemma 1.5.

$$\mathbb{E}[\ell(\hat{h}(X), Y)] - \mathbb{E}[\ell(h^*(X), Y))] \leq 2\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\hat{\beta}(y) - \beta(y)\|_2\right] \mathbb{E}[\|\phi(X)\|_2]. \tag{13}$$

Proof. By applying the definition of $R(y|x) := \mathbb{E}[\ell(y, Y)|x]$ and iterated expectation, the LHS of (13) can be written as

$$\mathbb{E}[\ell(\hat{h}(X), Y)] - \mathbb{E}[\ell(h^*(X), Y))] = \mathbb{E}[R(\hat{h}(X)|X) - R(h^*(X)|X)].$$

We will now bound $R(\hat{h}(x)|x) - R(h^*(x)|x)$ as follows. Observe that for every $x \in \mathcal{X}$, $\hat{h}(x)$ is a minimizer of the auxiliary problem, $\hat{h}(x) \in \arg\min_{y\in\mathcal{Y}} \hat{R}(y|x)$. As a shorthand notation, we will define $\hat{r}_x(g) := \hat{R}(g(x)|x)$ and $r_x(g) := R(g(x)|x)$ for any predictor g. Then,

$$\begin{aligned} R(\hat{h}(x)|x) - R(h^*(x)|x) &= r_x(\hat{h}) - r_x(h^*) \\ &= r_x(\hat{h}) - \hat{r}_x(\hat{h}) + \hat{r}_x(\hat{h}) - \hat{r}_x(h^*) + \hat{r}_x(h^*) - r_x(h^*) \\ &\leq r_x(\hat{h}) - \hat{r}_x(\hat{h}) + \hat{r}_x(h^*) - r_x(h^*), \end{aligned}$$

where the inequality is due to $\hat{r}_x(\hat{h}) - \hat{r}_x(h^*) \leq 0$, which follows from $\hat{h}(x)$ being a minimizer of $\hat{R}(y|x)$. Now note that $|r_x(\hat{h}) - \hat{r}_x(\hat{h})| = |R(\hat{h}(x)|x) - \hat{R}(\hat{h}(x)|x)| \leq \sup_{y\in\mathcal{Y}} |R(y|x) - \hat{R}(y|x)|$. Likewise, we can apply the same bound to $|\hat{r}_x(h^*) - r_x(h^*)|$. Therefore,

$$\begin{aligned} R(\hat{h}(x)|x) - R(h^*(x)|x) &\leq |r_x(\hat{h}) - \hat{r}_x(\hat{h})| + |\hat{r}_x(h^*) - r_x(h^*)| \\ &\leq 2\sup_{y\in\mathcal{Y}} |R(y|x) - \hat{R}(y|x)|. \end{aligned}$$

To complete the proof, substitute $R(y|x) = \beta(y)^T\phi(x)$ and $\hat{R}(y|x) = \hat{\beta}(y)^T\phi(x)$ into the last line, then apply Cauchy–Schwarz inequality, $|\beta(y)^T\phi(x) - \hat{\beta}(y)^T\phi(x)| \leq \|\beta(y) - \hat{\beta}(y)\|_2\|\phi(x)\|_2$. □

1.6.2 Proof of Theorem 5.1: Step 2

Lemma 1.6.

$$\sup_{y\in\mathcal{Y}} \|\hat{\beta}(y) - \beta(y)\|_2^2 \leq \sup_{y\in\mathcal{Y}} \frac{\|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2^2}{\nu m}$$

Proof. For a fixed y, $\frac{1}{m}\|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2^2$ can be understood as the mean squared error (MSE) of estimating $\Phi\beta(y)$ with $\Phi\hat{\beta}(y)$, which can be expressed as

$$\frac{1}{m}\|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2^2 = (\hat{\beta}(y) - \beta(y))^T \left(\frac{1}{m}\Phi^T\Phi\right)(\hat{\beta}(y) - \beta(y)).$$

Let $\nu_1 \leq \nu_2 \leq \ldots \leq \nu_q$ be the eigenvalues of $\frac{1}{m}\Phi^T\Phi$. Then, it follows that $u^T(\frac{1}{m}\Phi^T\Phi)u \geq \nu_1\|u\|_2^2$ for any $u \in \mathbb{R}^q$. Applying this to the above with $u := \hat{\beta}(y) - \beta(y)$ and $\nu := \nu_1$, we obtain

$$\nu\|\hat{\beta}(y) - \beta(y)\|_2^2 \leq \frac{1}{m}\|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2^2.$$

Taking the supremum over $y \in \mathcal{Y}$ on both sides completes the proof. □

We will now bound the MSE, $\frac{1}{m}\|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2^2$, using a standard approach for analyzing linear regression. In our setting, the design matrix is Φ and the vector of observed dependent variables are $[\ell(y, y^{(i)}) \ \cdots \ \ell(y, y^{(m)})]^T$. The underlying linear regression model is as follows,

$$\ell(y, Y) = \beta(y)^T \phi(X) + \varepsilon(y, X). \tag{14}$$

Here $\varepsilon(y, X)$ is a zero-mean random noise that can depend on X. In what follows, we will derive the results for any fixed y. Let $\varepsilon_i(y) := \varepsilon(y, x^{(i)})$ and $\varepsilon(y) := [\varepsilon_1(y) \ \cdots \ \varepsilon_m(y)]^T$. The following lemma shows how $\|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2$ can be bounded in terms of $\varepsilon(y)$ and Φ.

Proposition 1.7. *Suppose $u_1, \ldots, u_r \in \mathbb{R}^m$, $r \leq q$ is a set of orthornormal basis vectors for the column space of Φ, and $\tilde{\varepsilon}(y) := [\varepsilon(y)^T u_1 \ \cdots \ \varepsilon(y)^T u_r]^T$. Then,*

$$\|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2^2 \leq 4\|\tilde{\varepsilon}(y)\|_2^2.$$

Proof. This follows from a standard analysis of linear regression. The proof is available in the literature, e.g., in [5]. □

1.6.3 Proof of Theorem 5.1: Step 3

Combining the two lemmas in Step 1 and 2, we now have

$$\mathbb{E}[\ell(\hat{h}(X), Y)] - \mathbb{E}[\ell(h^*(X), Y))] \leq \frac{2}{\sqrt{\nu m}}\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2\right] \mathbb{E}[\|\phi(X)\|_2].$$

What remains is to bound $\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2\right]$. To that end, we will strengthen the result in Proposition 1.7 with supremum over $y \in \mathcal{Y}$. A key step is the following decomposition,

$$\begin{aligned}
\sup_{y\in\mathcal{Y}} \|\tilde{\varepsilon}(y)\|_2^2 &= \sup_{y\in\mathcal{Y}} \sum_{j=1}^{r} \tilde{\varepsilon}_j(y)^2 \\
&\leq \sum_{j=1}^{r} \sup_{y\in\mathcal{Y}} \left(\tilde{\varepsilon}_j(y)^2 - \mathbb{E}[\tilde{\varepsilon}_j(y)^2] + \mathbb{E}[\tilde{\varepsilon}_j(y)^2]\right) \\
&\leq \sum_{j=1}^{r} \left(\sup_{y\in\mathcal{Y}} \left(\tilde{\varepsilon}_j(y)^2 - \mathbb{E}[\tilde{\varepsilon}_j(y)^2]\right) + \sup_{y\in\mathcal{Y}} \mathbb{E}[\tilde{\varepsilon}_j(y)^2]\right).
\end{aligned} \tag{15}$$

We now derive some properties of $\tilde{\varepsilon}_j(y)^2 - \mathbb{E}[\tilde{\varepsilon}_j(y)^2]$ and $\tilde{\varepsilon}_j(y)$ that will allow us to bound their expectations over the supremum of y. Recall that $\tilde{\varepsilon}_j(y) = \varepsilon(y)^T u_j$ and $\varepsilon_i(y) := \varepsilon(y, x^{(i)})$. By the assumption of the theorem, $\ell(y, Y)$ is bounded above by L w.p. 1. This implies that $\varepsilon_i(y)$ can only vary in an interval of length L and thus have variance no more than $L^2/4$. More generally, we can characterize these random variables as either *sub-Gaussian* or *sub-Exponential*, defined as follows.

Definition (Sub-Gaussian) A r.v. Z is said to be sub-Gaussian with variance proxy σ^2 if $\mathbb{E}[Z] = 0$ and $\mathbb{E}[e^{sZ}] \leq e^{s^2\sigma^2/2}$, for any $s \in \mathbb{R}$.

Definition (Sub-Exponential) A r.v. Z is said to be sub-Exponential with parameter λ if $\mathbb{E}[Z] = 0$ and $\mathbb{E}[e^{sZ}] \leq e^{s^2\lambda^2/2}$, for any $s \in [-1/\lambda, 1/\lambda]$.

By Hoeffding's lemma, the boundedness of $\varepsilon_i(y)$ implies that it is sub-Gaussian with variance proxy $L^2/4$. The following proposition states several implications of this fact.

Proposition 1.8.

1. *$\tilde{\varepsilon}_j(y)$ is sub-Gaussian with variance proxy $L^2/4$.*
2. *$\tilde{\varepsilon}_j(y)^2 - \mathbb{E}[\tilde{\varepsilon}_j(y)^2]$ is sub-Exponential with parameter $4L^2$.*
3. *$\mathbb{E}[\tilde{\varepsilon}_j^2(y)] \leq L^2$.*

Proof. These follow from known properties of sub-Gaussian and sub-Exponential r.v's. To prove property 1, note that $\varepsilon(y)$ is a vector of i.i.d. sub-Gaussian r.v.'s with variance proxy $L^2/4$ and u_j is a basis vector with $\|u_j\|_2 = 1$. These conditions imply that their inner product, $\tilde{\varepsilon}_j(y) = \varepsilon(y)^T u_j$ is also sub-Gaussian with the same parameter (see [5] for a proof). Property 2 follows from the fact that if any r.v. Z is sub-Gaussian with variance proxy σ^2, then $Z^2 - \mathbb{E}[Z^2]$ is sub-Exponential with parameter $16\sigma^2$ (see the Appendix of [6] for a proof). A proof of property 3 is available in [5]. □

Property 3 of Proposition 1.8 immediately implies that $\sup_{y\in\mathcal{Y}} \mathbb{E}[\tilde{\varepsilon}_j(y)^2] \le L^2$. Our goal now is to use property 1 and 2 to obtain a maximal inequality for $\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \left(\tilde{\varepsilon}_j(y)^2 - \mathbb{E}[\tilde{\varepsilon}_j(y)^2]\right)\right]$. For convenience, let $\gamma_j(y) := \tilde{\varepsilon}_j(y)^2 - \mathbb{E}[\tilde{\varepsilon}_j(y)^2]$ and $\mathcal{Y} := \{y_1, \ldots, y_{|\mathcal{Y}|}\}$. For any $s \in (0, 1/4L^2]$,

$$\begin{aligned}
\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \gamma_j(y)\right] &= \mathbb{E}\left[\max_{k=1,\ldots,|\mathcal{Y}|} \gamma_j(y_k)\right] \\
&= \frac{1}{s}\mathbb{E}\left[\log \exp\left(s \max_{k=1,\ldots,|\mathcal{Y}|} \gamma_j(y_k)\right)\right] \\
&\le \frac{1}{s}\log \mathbb{E}\left[\exp\left(s \max_{k=1,\ldots,|\mathcal{Y}|} \gamma_j(y_k)\right)\right] \\
&= \frac{1}{s}\log \mathbb{E}\left[\max_{k=1,\ldots,|\mathcal{Y}|} \exp(s\gamma_j(y_k))\right] \\
&\le \frac{1}{s}\log \sum_{k=1}^{|\mathcal{Y}|} \mathbb{E}\left[\exp(s\gamma_j(y_k))\right] \\
&\le \frac{1}{s}\log \sum_{k=1}^{|\mathcal{Y}|} \exp(8s^2L^4) \\
&= \frac{\log|\mathcal{Y}|}{s} + 8sL^4.
\end{aligned}$$

The first inequality follows from Jensen's inequality and the third inequality is due to Property 2 of Proposition 1.8. By letting $s = 1/4L^2$, the last line becomes $2L^2(2\log|\mathcal{Y}| + 1)$. Putting all these together in (15),

$$\begin{aligned}
\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\tilde{\varepsilon}(y)\|_2^2\right] &\le \sum_{j=1}^{r} \left(2L^2(2\log|\mathcal{Y}| + 1) + L^2\right) \\
&= rL^2(4\log|\mathcal{Y}| + 3)
\end{aligned}$$

Finally, we apply the following inequalities to complete the proof,

$$\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2\right] \le \sqrt{\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\Phi\hat{\beta}(y) - \Phi\beta(y)\|_2^2\right]} \le 2\sqrt{\mathbb{E}\left[\sup_{y\in\mathcal{Y}} \|\tilde{\varepsilon}(y)\|_2^2\right]}.$$

1.7 Accounting for intercept in Proposition 2.1

We briefly discuss how to account for intercept in (2). In linear regression, this is often done by centering both the dependent and output variables at their respective empirical means, and then solving the least squares problem without an intercept. The output is then translated by the mean to obtain a final estimate, which can be shown to be equivalent to the output of a model with intercept [7]. However, if we apply the same approach to (2), the result may not be what we expect: because the kernel implicitly defines a mapping $\phi(x) \mapsto \mathcal{H}$ from the original input space to a feature space, which can be highly nonlinear, a point that is centered in $\mathcal{X}$ need not be centered in $\mathcal{H}$. This can result in biases.

One alternative is to center the inputs in the feature space by working only with the inner products. Let us define $\tilde{\phi}_i := \phi_i - \bar{\phi}$ to be a centered input in the feature space, where $\phi_i := \phi(x^{(i)})$ and $\bar{\phi} := \frac{1}{m}\sum_{i=1}^{m} \phi(x^{(i)})$. Then, the gram matrix in (3) can be characterized by

$$\begin{aligned}
\langle\tilde{\phi}_i, \tilde{\phi}_j\rangle_{\mathcal{H}} &= \langle\phi_i - \bar{\phi}, \phi_j - \bar{\phi}\rangle_{\mathcal{H}} \\
&= \langle\phi_i, \phi_j\rangle_{\mathcal{H}} - \langle\phi_i, \bar{\phi}\rangle_{\mathcal{H}} - \langle\phi_j, \bar{\phi}\rangle_{\mathcal{H}} + \langle\bar{\phi}, \bar{\phi}\rangle_{\mathcal{H}} \\
&:= \tilde{k}(x^{(i)}, x^{(j)}).
\end{aligned}$$

It can be shown that the resulting $\tilde{k}$ is also a positive definite kernel, and its gram matrix $\tilde{K}$ is a *centered kernel matrix* [8] that can be written as

$$\tilde{K} = \left(I - \frac{11^T}{m}\right) K \left(I - \frac{11^T}{m}\right),$$

where 1^T is a row vector of ones of appropriate dimension. The output variables can be centered as follows. For each y, let us define $\bar{\ell}_y := \frac{1}{m}\sum_{i=1}^{m} \ell(y, y^{(i)})$ to be the empirical mean of losses. We center the vector L_y defined in (3) by letting $\tilde{L}_y := L_y - \bar{\ell}_y 1$. Then, just as in linear regression, we estimate $\hat{R}(y|x)$ by solving (3)

with all the centered inputs, and then translate the result by $\bar{\ell}_y$,

$$\begin{aligned}\hat{R}(y|x) &= \tilde{L}_y^T(\tilde{K}+\lambda I)^{-1}\tilde{v}(x) + \bar{\ell}_y \\ &= \sum_{i=1}^{m}(\tilde{w}_i(x) - u(x))\ell(y, y^{(i)}) + \frac{1}{m}\sum_{i=1}^{m}\ell(y, y^{(i)}).\end{aligned}$$

Here $\tilde{w}(x) := (\tilde{K}+\lambda I)^{-1}\tilde{v}(x)$, $\tilde{v}(x) := [\tilde{k}(x, x^{(i)})]_{i=1}^m$ and $u(x) := \frac{1}{m}\sum_{i=1}^m \tilde{w}_i(x)$. In the above expression, we can interpret the second term $\frac{1}{m}\sum_{i=1}^m \ell(y, y^{(i)})$ as a baseline sample average approximation of the risk function, while the first term as the correction after observing x.

2 Appendix B: Supplemental experimental materials

2.1 Summary of data sets

Table 1: A summary of the attributes of each data set. Cardinality is the average number of labels per sample. Max Depth and Avg Depth correspond to the depth of the hierarchy and the average depth of the labels, respectively.

Data set	#Features	#Labels	#Train	#Test	Cardinality	Max Depth	Avg Depth
enron	1001	56	988	660	5.30	2	1.18
rcv1v2	47236	103	3000	3000	3.23	3	1.40
imclef07a	80	96	2000	1006	3.00	2	1.57
imclef07d	80	46	2000	1006	3.00	2	1.48
pheno_go	276	296	1005	581	5.44	10	4.01
spo_go	84	508	2434	1263	5.58	10	4.24
gasch1_go	173	513	2477	1281	5.57	10	4.24
gasch2_go	52	513	2485	1288	5.56	10	4.24

The data sets are available at these sources:

1. `http://kt.ijs.si/DragiKocev/PhD/resources/doku.php?id=hmc_classification`
2. `https://dtai.cs.kuleuven.be/clus/hmcdatasets`
3. `https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/multilabel.html` – rcv1v2 (topics; subset1)[7].

For imclef07a and imclef07d, we only used the first 2000 training samples out of the 10000. In the last 4 data sets, we combined the training and validation set into a single training set for cross-validated parameter tuning. Furthermore, only the first connected component of the full hierarchy (>4000 labels) is used, which is then trimmed down by discarding labels with less than 3 positive instances.

2.2 PR Curves

The precision-recall curves are shown in Figure 1. For LS-ECRM, the curve is obtained by varying the relative costs of false positives and false negatives by modifying the Hamming loss,

$$\ell(y, y') := \frac{1}{d}\sum_{j=1}^{d}\left(C_1\mathbb{1}(y_j = 0, y'_j = 1) + C_2\mathbb{1}(y_j = 1, y'_j = 0)\right),$$

for some penalty factor C_1 and C_2. This can be transformed into a linear objective in (10). For BR-SVM, H-SVM and Clus-HMC, it is done by varying the classification thresholds.

[7]The topic hierarchy is available at `http://www.jmlr.org/papers/volume5/lewis04a/lyrl2004_rcv1v2_README.htm`

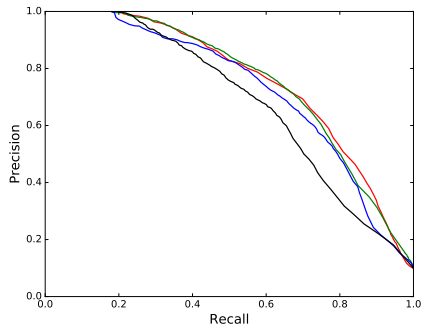

(c) enron

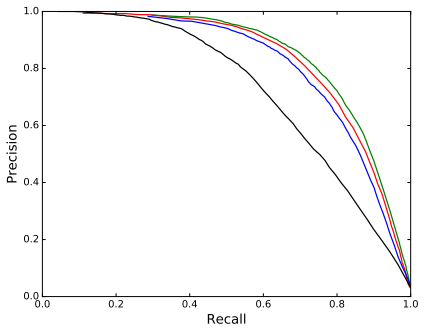

(d) rcv1v2

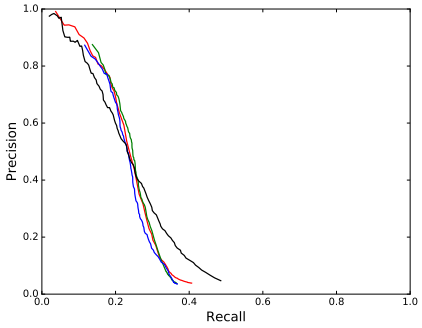

(e) imclef07a

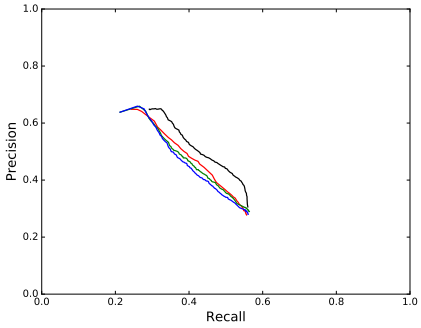

(f) imclef07d

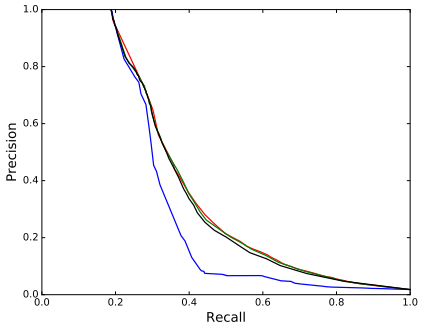

(g) pheno_go

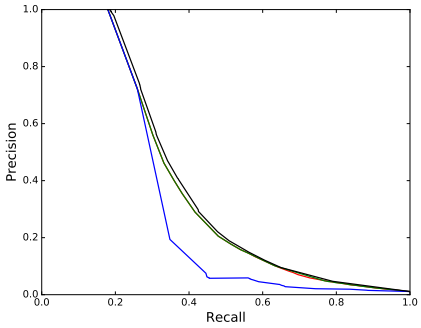

(h) spo_go

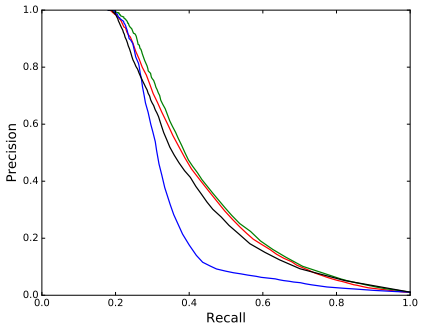

(i) gasch1_go

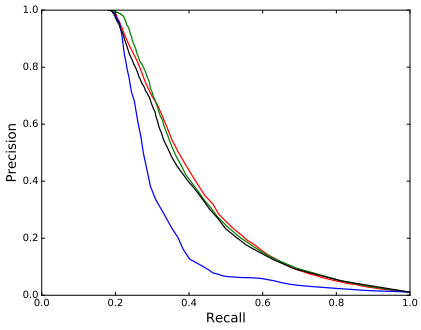

(j) gasch2_go

Figure 1: Precision-recall curves of LS-ECRM (red), BR-SVM (green), H-SVM (blue) and Clus-HMC (black) on various data sets.

References

[1] Schölkopf, B., Herbrich, R., & Smola, A. J. (2001). A generalized representer theorem. In Computational learning theory (pp. 416-426). Springer Berlin Heidelberg.

[2] Steiglitz, K., & Papadimitriou, C. H. (1982). Combinatorial optimization: Algorithms and complexity. Printice-Hall, New Jersey.

[3] Bertsimas, D., & Tsitsiklis, J. N. (1997). Introduction to linear optimization (Vol. 6, pp. 65-67). Belmont, MA: Athena Scientific.

[4] Conforti, M., Cornuéjols, G., & Zambelli, G. (2014). Integer programming (Vol. 271). Berlin: Springer.

[5] Rigollet, P. (2015) 18.S997: High Dimensional Statistics Lecture Notes. Retrieved from `http://www-math.mit.edu/~rigollet/PDFs/RigNotes15.pdf`.

[6] Honorio, J., & Jaakkola, T. (2014). Tight bounds for the expected risk of linear classifiers and PAC-Bayes finite-sample guarantees.

[7] Hastie, T., Tibshirani, R., Friedman, J., & Franklin, J. (2005). The elements of statistical learning: data mining, inference and prediction. The Mathematical Intelligencer, 27(2), 83-85.

[8] Cortes, C., Mohri, M., & Rostamizadeh, A. (2012). Algorithms for learning kernels based on centered alignment. The Journal of Machine Learning Research, 13(1), 795-828.